

64 bits

[Andrey Karpov](#)

OOO "Program Verification Systems"

May 2010

[Abstract](#)

[Introduction](#)

[The history of 64-bit systems](#)

[Applied programming and 64-bit systems](#)

[Intel 64 \(AMD64\) architecture](#)

[64-bit operating systems](#)

[WoW64](#)

[Win64 program model](#)

[Address space](#)

[64-bit software development](#)

[References](#)

Abstract

The article reveals the meaning of the term "64 bits". It briefly discusses the history of 64-bit system development, describes the most popular 64-bit processors of the Intel 64 architecture and the 64-bit Windows operating system.

Introduction

By the term "64-bit", within the scope of computer architecture, we understand 64-bit integers and other data types with the size 64 bits. By "64-bit systems" 64-bit microprocessor architectures (for example, EM64T, IA-64) or 64-bit operating systems (for example, Windows XP Professional x64 Edition) may be understood. There are also compilers that generate 64-bit program code.

In this article, we will discuss various aspects related to 64-bit technologies. It is intended for programmers who want to start developing 64-bit programs and is oriented on Windows-developers since the task of studying 64-bit systems is the most relevant for them.

The history of 64-bit systems

64 bits have just recently come into life of most users and applied programmers. But the history of working with 64-bit data is rather long.

1961: IBM delivers the IBM 7030 Stretch supercomputer, which uses 64-bit data words and 32- or 64-bit instruction words.

1974: Control Data Corporation launches the CDC Star-100 vector supercomputer, which uses a 64-bit word architecture (previous CDC systems were based on a 60-bit architecture).

1976: Cray Research delivers the first Cray-1 supercomputer, which is based on a 64-bit word architecture and will form the basis for later Cray vector supercomputers.

1985: Cray releases UNICOS, the first 64-bit implementation of the Unix operating system.

1991: MIPS Technologies produces the first 64-bit microprocessor, the R4000, which implements the MIPS III ISA, the third revision of their MIPS architecture. The CPU is used in SGI graphics workstations starting with the IRIS Crimson. Kendall Square Research deliver their first KSR1 supercomputer, based on a proprietary 64-bit RISC processor architecture running OSF/1.

1992: Digital Equipment Corporation (DEC) introduces the pure 64-bit Alpha architecture which was born from the PRISM project.

1993: DEC releases the 64-bit DEC OSF/1 AXP Unix-like operating system (later renamed Tru64 UNIX) for its systems based on the Alpha architecture.

1994: Intel announces plans for the 64-bit IA-64 architecture (jointly developed with Hewlett-Packard) as a successor to its 32-bit IA-32 processors. A 1998 to 1999 launch date is targeted. SGI releases IRIX 6.0, with 64-bit support for the R8000 chip set.

1995: Sun launches a 64-bit SPARC processor, the UltraSPARC. Fujitsu-owned HAL Computer Systems launches workstations based on a 64-bit CPU, HAL's independently designed first-generation SPARC64. IBM releases the A10 and A30 microprocessors, 64-bit PowerPC AS processors. IBM also releases a 64-bit AS/400 system upgrade, which can convert the operating system, database and applications.

1996: Nintendo introduces the Nintendo 64 video game console, built around a low-cost variant of the MIPS R4000. HP releases an implementation of the 64-bit 2.0 version of their PA-RISC processor architecture, the PA-8000.

1997: IBM releases the RS64 line of 64-bit PowerPC/PowerPC AS processors.

1998: Sun releases Solaris 7, with full 64-bit UltraSPARC support.

1999: Intel releases the instruction set for the IA-64 architecture. AMD publicly discloses its set of 64-bit extensions to IA-32, called x86-64 (later branded AMD64).

2000: IBM ships its first 64-bit ESA/390-compatible mainframe, the zSeries z900, and its new z/OS operating system.

2001: Intel finally ships its 64-bit processor line, now branded Itanium, targeting high-end servers. It fails to meet expectations due to the repeated delays in getting IA-64 to market. NetBSD is the first operating system to run on the Intel Itanium processor at the processor's release. Additionally, Microsoft releases Windows XP 64-Bit Edition, also for the Itanium's IA-64 architecture, although it was capable of running 32-bit applications through an execution layer WoW64.

2003: AMD introduces its Opteron and Athlon 64 processor lines, based on its AMD64 architecture which is the first x86 based 64 bit processor architecture. Apple also ships the 64-bit "G5" PowerPC 970 CPU courtesy of IBM. Intel maintains that its Itanium chips would remain its only 64-bit processors.

2004: Intel, reacting to the market success of AMD, admits it has been developing a clone of the AMD64 extensions named IA-32e (later renamed EM64T, then yet again renamed to Intel 64). Intel also ships updated versions of its Xeon and Pentium 4 processor families supporting the new instructions.

2004: VIA Technologies announces the Isaiah 64-bit processor.

2005: On January 31, Sun releases Solaris 10 with support for AMD64 / Intel 64 processors. On April 30, Microsoft releases Windows XP Professional x64 Edition for AMD64 / Intel 64 processors.

2006: Sony, IBM, and Toshiba begin manufacturing of the 64-bit Cell processor for use in the

PlayStation 3, servers, workstations, and other appliances. Microsoft releases Windows Vista, including a 64-bit version for AMD64 / Intel 64 processors that retains 32-bit compatibility. All Windows applications and components are 64-bit, although many also have their 32-bit versions included for compatibility with plugins.

2009: Microsoft's Windows 7, like Windows Vista, includes a full 64-bit version for AMD64/Intel 64 processors and most new computers are loaded by default with a 64-bit version. Apple's Mac OS X 10.6, "Snow Leopard" is shipped with a 64-bit kernel for AMD64 / Intel 64 processors, although only certain recent models of Apple computers will run this by default. Most applications bundled with Mac OS X 10.6 are now also 64-bit.

To learn more about the history of 64-bit system development, see the article "[The Long Road to 64 Bits](#)" [1] by John Mashey and in the Wikipedia-entry "[64-bit](#)" [2].

Applied programming and 64-bit systems

By the moment of writing this article (2010), there are two most popular and widely used [64-bit](#) microprocessor architectures: IA64 and Intel 64.

1. **IA-64** is a 64-bit microprocessor architecture developed by Intel and Hewlett Packard companies jointly. It is implemented in Itanium and Itanium 2 microprocessors. To learn more about the architecture IA-64, see the Wikipedia entry: "[Itanium](#)". The Itanium architecture is supported by most server vendors: Bull, Fujitsu, Fujitsu Siemens Computers, Hitachi, HP, NEC, SGI and Unisys. These manufacturers joined Intel and many software developers to create [Itanium Solutions Alliance](#) with the purpose of promoting the architecture and speeding up the software porting rate.
2. **Intel 64** (AMD64 / x86-64 / x64 / EM64T) is an extension of the [x86](#) architecture with complete backward compatibility. There are a lot of names for this architecture and it causes some confusion, although all these names actually mean the same: x86-64, AA-64, Hammer Architecture, AMD64, Yamhill Technology, EM64T, IA-32e, Intel 64, x64. To learn more about how so many names appeared, see the Wikipedia entry: "[X86-64](#)". Processors with the Intel 64 architecture are widely used in personal computers. It is most likely that your computer has a processor with this very architecture.

You should understand that IA-64 and Intel 64 (AMD64) are absolutely different microprocessor architectures which are absolutely incompatible with each other. Further in this article, we will discuss only the Intel 64 (x64 / AMD64) architecture as the most popular among Windows applied software developers. For short, the Intel 64 program model available in the 64-bit Windows is called [Win64](#).

Intel 64 (AMD64) architecture

The architecture [Intel 64](#) we are discussing here, is a simple yet powerful backward compatible extension of the obsolete industry architecture x86. It adds the 64-bit address space and extends register resources for the purpose of higher performance of recompiled 64-bit applications. The architecture supports obsolete 16-bit and 32-bit code of applications and operating systems without modifying or recompiling them.

A special feature of Intel 64 is support of sixteen 64-bit general-purpose registers (there were eight 32-bit registers in x86-32). There is also support for 64-bit arithmetic and logical operations over integer numbers. To address new registers, instructions were provided with "register extension prefixes" for which the range of codes 40h-4Fh was chosen which are used for INC <register> and DEC <register> instructions in 32-bit and 16-bit modes. The instructions INC and DEC must be coded in a more general, two-byte form in the 64-bit mode.

Registers:

16 integer 64-bit general-purpose registers (RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, R8 — R15),

8 80-bit floating-point registers (ST0 — ST7),

8 64-bit Multimedia Extensions registers (MM0 — MM7, they share space with the registers ST0 — ST7),

16 128-bit SSE registers (XMM0 — XMM15),

64-bit RIP pointer and 64-bit flag register RFLAGS.

The need for a 64-bit architecture is determined by applications that need a large address space. First of all, these are high-performance servers, data managers, CAD and, of course, games. These applications will get significant benefits from the 64-bit address space and extended number of registers. Few registers available in the obsolete x86 architecture limit performance in computational tasks. The increased number of registers provides sufficient performance for many applications.

Let us point out the main advantages of the x86-64 architecture:

- 64-bit address space;
- extended register set;
- instruction set familiar to developers;
- capability to launch obsolete 32-bit applications in a 64-bit operating system;
- capability to use 32-bit operating systems.

64-bit operating systems

Nearly all contemporary operating systems nowadays have versions for the Intel 64 architecture. For example, Microsoft ships Windows XP x64. Largest UNIX-system developers also ship 64-bit versions, for example, Linux Debian 3.5 x86-64. But it does not mean that the whole code of such a system is completely 64-bit. Some code of the operating system and many applications might well remain 32-bit since Intel 64 provides backward compatibility with 32-bit applications. For example, the 64-bit Windows version employs a special mode [WoW64](#) (Windows-on-Windows 64) which translates calls of 32-bit applications to the resources of the 64-bit operating system.

Further in the article, we will discuss only 64-bit operating systems of the Windows family.

WoW64

Windows-on-Windows 64-bit (WoW64) is a subsystem of the Windows operating system that allows you to launch 32-bit applications under all 64-bit Windows versions.

The WoW64 subsystem does not support the following programs:

- programs compiled for 16-bit operating systems;
- kernel mode programs compiled for 32-bit operating systems.

There are some differences in various implementations of WoW64 determined by the processor architecture. For example, the 64-bit Windows version developed for the Intel Itanium 2 processor uses WoW64 to emulate [x86](#) instructions. Such emulation is very resource-intensive in comparison with WoW64 implemented for the [Intel 64](#) architecture since the system needs to switch to the compatibility

mode from the 64-bit mode when executing 32-bit programs.

WoW64 under the Intel 64 (AMD64 / x64) architecture does not require emulation of instructions. The WoW64 subsystem here emulates only the 32-bit environment through an additional layer between a 32-bit application and the 64-bit Windows API. This layer is thin in some places and thicker in some others. For an average program, the performance loss determined by this layer will be about 2%. This value might be a bit higher for some programs. Two percent is very little but you should take into account that 32-bit applications work much slower under the 64-bit Windows operating system than in their native 32-bit environment.

Compilation of 64-bit code does not only eliminate the necessity of using WoW64 but also provides some additional performance gain. It is determined by architectural changes in the microprocessor, such as an increased number of general-purpose registers. For an average program, you may expect a 5-15% performance gain after a mere recompilation.

Because of the WoW64 layer, 32-bit programs work less efficiently in the 64-bit environment than in the 32-bit one. But still simple 32-bit applications may get one benefit of being launched in the 64-bit environment. Perhaps you know that a program built with the switch `/LARGEADDRESSAWARE:YES` can allocate up to 3 Gbytes of memory if a 32-bit Windows operating system is launched with the `/3gb` switch. Well, the same 32-bit program launched on a 64-bit system can allocate about 4 Gbytes of memory (in practice it is about 3.5 GBytes).

The WoW64 subsystem isolates 32-bit programs from 64-bit ones by redirecting the calls to files and register. It prevents an accidental access of 32-bit programs to data of 64-bit ones. For example, a 32-bit application that launches a DLL file from the folder `%systemroot%\System32` might accidentally address a 64-bit DLL which is incompatible with this 32-bit program. To avoid this, the WoW64 subsystem redirects the access from the folder `%systemroot%\System32` into the folder `%systemroot%\SysWOW64`. This redirection allows to prevent compatibility errors since a DLL file is required which is created specially for 32-bit applications.

To learn more about the file system and register redirection mechanisms, see the MSDN section "[Running 32-bit Applications](#)".

Win64 program model

Like in [Win32](#), the size of pages in Win64 is 4 Kbytes. The first 64 Kbytes of the address space are never displayed, that is the least correct address is `0x10000`. Unlike Win32, system DLLs can be loaded for more than 4 Gbytes.

The specific feature of compilers for Intel 64 is that they can most effectively use registers to pass parameters into functions instead of using the stack. It allowed the developers of the Win64 architecture to get rid of such a notion as [calling convention](#). In Win32, you may use different conventions: `__stdcall`, `__cdecl`, `__fastcall`, etc. In Win64, there is only one calling convention. Let us consider an example of how four arguments of *integer* type are passed in registers:

- RCX: the first argument
- RDX: the second argument
- R8: the third argument
- R9: the fourth argument

The arguments after the first four *integers* are passed on the stack. To pass *float* arguments, XMM0-XMM3 registers are used as well as the stack.

The difference between the calling conventions results in an impossibility to use both 64-bit and 32-bit code in one program. In other words, if an application is compiled for the 64-bit mode, all the DLL libraries being used must be also 64-bit.

Passing of parameters through registers is one of the innovations that make 64-bit programs faster than 32-bit ones. An additional performance gain might be obtained if you use 64-bit data types.

Address space

Although a 64-bit processor can theoretically address 16 Ebytes of memory (2^{64}), Win64 currently supports only 16 Tbytes (2^{44}). There are several reasons for that. Contemporary processors can provide an access only to 1 Tbyte (2^{40}) of physical memory. The architecture (but not the hardware part) can extend this space up to 4 Pbytes (2^{52}). But in this case you will need huge memory amounts for page tables that represent the memory.

Besides the mentioned restrictions, the amount of memory available in a particular 64-bit version of Windows operating system depends also upon the commercial policy of Microsoft company. Below is the information about how much memory is supported in different 64-bit Windows versions:

Windows XP Professional - 128 Gbyte;

Windows Server 2003, Standard - 32 Gbyte;

Windows Server 2003, Enterprise - 1 Tbyte;

Windows Server 2003, Datacenter - 1 Tbyte;

Windows Server 2008, Datacenter - 2 Tbyte;

Windows Server 2008, Enterprise - 2 Tbyte;

Windows Server 2008, Standard - 32 Gbyte;

Windows Server 2008, Web Server - 32 Gbyte;

Vista Home Basic - 8 Gbyte;

Vista Home Premium - 16 Gbyte;

Vista Business - 128 Gbyte;

Vista Enterprise - 128 Gbyte;

Vista Ultimate - 128 Gbyte;

Windows 7 Home Basic - 8 Gbyte;

Windows 7 Home Premium - 16 Gbyte;

Windows 7 Professional - 192 Gbyte;

Windows 7 Enterprise - 192 Gbyte;

Windows 7 Ultimate - 192 Gbyte;

64-bit software development

The issues of 64-bit software development are covered most thoroughly in the course "[Lessons on development of 64-bit C/C++ applications](#)". Contents:

- Lesson 01. [What 64-bit systems are.](#)
- Lesson 02. [Support of 32-bit applications.](#)
- Lesson 03. [Porting code to 64-bit systems. The pros and cons.](#)
- Lesson 04. [Creating the 64-bit configuration.](#)
- Lesson 05. [Building a 64-bit application.](#)
- Lesson 06. [Errors in 64-bit code.](#)
- Lesson 07. [The issues of detecting 64-bit errors.](#)
- Lesson 08. [Static analysis for detecting 64-bit errors.](#)
- Lesson 09. [Pattern 01. Magic numbers.](#)
- Lesson 10. [Pattern 02. Functions with variable number of arguments.](#)
- Lesson 11. [Pattern 03. Shift operations.](#)
- Lesson 12. [Pattern 04. Virtual functions.](#)
- Lesson 13. [Pattern 05. Address arithmetic.](#)
- Lesson 14. [Pattern 06. Changing an array's type.](#)
- Lesson 15. [Pattern 07. Pointer packing.](#)
- Lesson 16. [Pattern 08. Memsizes in unions.](#)
- Lesson 17. [Pattern 09. Mixed arithmetic.](#)
- Lesson 18. [Pattern 10. Storage of integer values in double.](#)
- Lesson 19. [Pattern 11. Serialization and data interchange.](#)
- Lesson 20. [Pattern 12. Exceptions.](#)
- Lesson 21. [Pattern 13. Data alignment.](#)
- Lesson 22. [Pattern 14. Overloaded functions.](#)
- Lesson 23. [Pattern 15. Growth of structures' sizes.](#)

Lesson 24. [Phantom errors](#).

Lesson 25. [Working with patterns of 64-bit errors in practice](#).

Lesson 26. [Optimization of 64-bit programs](#).

Lesson 27. [Peculiarities of creating installers for a 64-bit environment](#).

Lesson 28. [Estimating the cost of 64-bit migration of C/C++ applications](#).

The authors of the course are the workers of the company "Program Verification Systems" that develop the static code analyzer Viva64 intended for detecting errors in 64-bit programs. On the company's site, you may find many other resources on 64-bit application development and migration of 32-bit applications to 64-bit systems. As an example, see the [article review section](#) on the topic related to 64-bit technologies.

References

1. John R. Mashey. The Long Road to 64 Bits. <http://www.viva64.com/go.php?url=321>
2. Wikipedia. 64-bit. <http://www.viva64.com/go.php?url=203>